

MULTICOMPONENT SUBSPACE CHIRP PARAMETER ESTIMATION USING DISCRETE FRACTIONAL FOURIER ANALYSIS

Daniel Jackson Peacock* and Balasubramaniam Santhanam†
Department of E.C.E.
University of New Mexico
Albuquerque, NM, USA: 87131
jpeacock,bsanathan@unm.edu

ABSTRACT

The Discrete Fractional Fourier Transform is a useful tool for multicomponent chirp parameter estimation, however much of the current work ignores the presence of noise. In recent work, a projection-subspace approach to DFrFT-based multicomponent chirp parameter estimation was proposed, to exploit the robustness of subspace decomposition methods. This paper refines the prior projection-subspace work to overcome limitations caused by the projection preprocessing, and presents a quantitative analysis of its performance as compared to the Cramér-Rao lower bound.

KEY WORDS

Discrete Fractional Fourier Transform, multicomponent chirp parameter estimation, subspace decomposition

1 Introduction

Chirp signals are sinusoidal waveforms with linearly changing instantaneous frequency. They find wide applications in radar systems, including synthetic aperture radar, biomedical applications [1] and can be used as a simplified model for bat echo-location signals. A robust method of multicomponent parameter estimation would enable the estimation of the vibrational frequency of a target and improve estimation performance in the presence of clutter.

The Discrete Fractional Fourier Transform (DFrFT) shows promise in multicomponent chirp parameter estimation as it generates a strong peak for each chirp whose location in the the 2D transform plane corresponds to the specific center frequency and chirp rate. The mapping between peak locations and chirp parameters was investigated in [2] resulting in a closed-form empirical approximation of the relationship, however this approximation introduces significant error in the parameter estimation.

Subspace decomposition techniques have also been investigated for use in conjunction with the DFrFT with

the aim of providing more robust and accurate estimation in the presence of noise [3]. The projection method used in this prior work, however, causes the peak to be suppressed, hindering the performance of the subspace decomposition. Finally, neither of the above papers investigated effects of noise on the performance of their proposed estimators.

In this paper, we will calculate the peak to parameter mapping at each point in the transform, greatly reducing the error caused by the approximate mapping. We investigate using different p-norms for the projection to better accentuate the peaks, and propose a cross-hairs method, where subspace decomposition is only performed on projections of thin slices centered on the peaks found by 2D peak detection. Finally, we quantify the performance of the estimators in the presence of noise using simulation data, and compare them to the Cramér-Rao lower bound.

2 Discrete Fractional Fourier Transform

The Fractional Fourier Transform (FrFT) is a generalization of the Fourier Transform. If time and frequency are treated as orthogonal axes, then the Fourier Transform is a 90° rotation in this plane, while the FrFT can generate signal representations at any angle of rotation in the plane [4]. The eigenvectors of the FrFT are Hermite-Gauss functions, which result in a kernel composed of chirps. Application of the FrFT to a chirp results in an impulse for a specific angle. The angle where this occurs corresponds to the chirp rate of the signal, and the location of the impulse in the transform corresponds to the center frequency, according to closed-form equations.

Discrete versions of the Fractional Fourier Transform have been developed by several researchers [5][6][7]. All the forms use approximations of the Hermite-Gauss functions for eigenvectors, which lead to approximate chirp functions as their kernel. The general form of the transform is given by [5]:

$$\mathbf{X}_a = \mathbf{V} \Lambda^{\frac{2a}{\pi}} \mathbf{V}^T \mathbf{x}, \quad (1)$$

where \mathbf{V} is a matrix of DFT eigenvectors and Λ is a diagonal matrix of DFT eigenvalues.

For this paper, we use the MA-CDFrFT as presented in [8], which is based on a centered DFT (whose kernel is

*This work was supported by Sandia National Laboratories, a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

†This work was supported by the US Department of Energy (award no. DE-FG52-08NA28782) and by the National Science Foundation (award no. IIS-0813747).

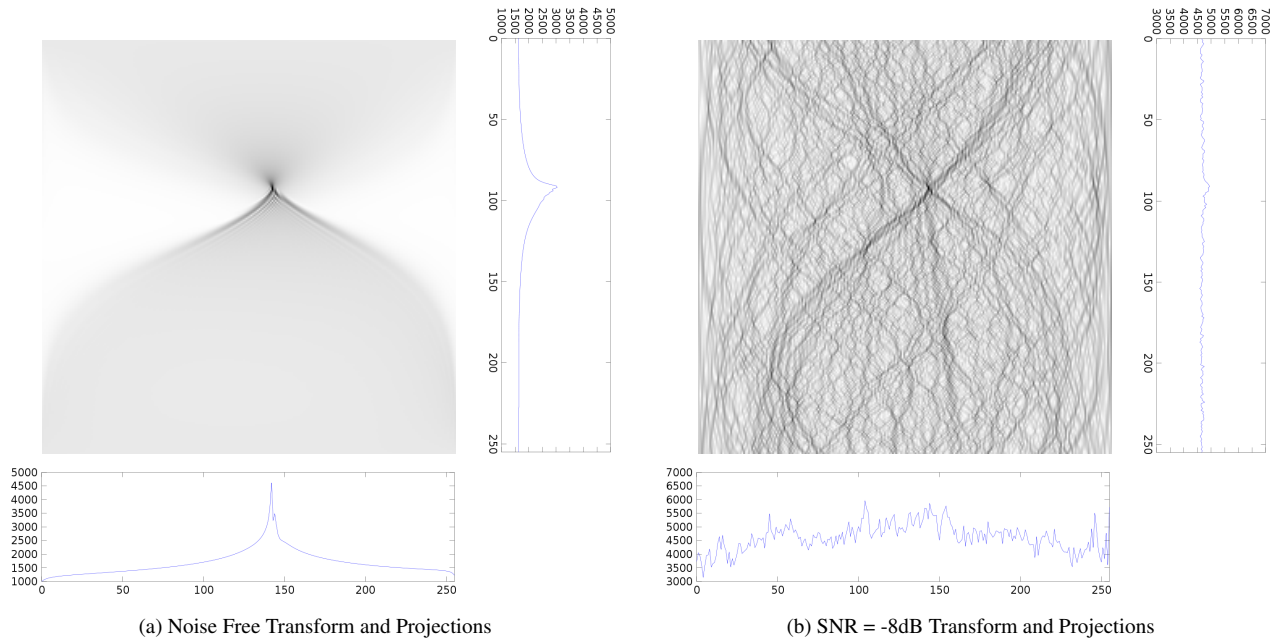


Figure 1: Effect of Noise on Projection: The MA-CDFrFT transform of a chimp with parameters $\alpha = \pi/2/N$, $\omega = \pi/8$ is shown above with and without noise added to the signal. Horizontal and vertical projections calculated using a 3-norm, show how the noise floor is higher in the 1D projections than in the 2D transform. The problem is more pronounced with 1-norm (absolute sum) projection.

generated by the Grünbaum tridiagonal commutator). When the length of the input is a multiple of four, symmetries of those eigenvectors allow for an efficient computation given by:

$$X_k[r] = \sum_{p=0}^{P-1} z_k[p] e^{-j \frac{2\pi}{P} pr}, \quad (2)$$

$$z_k[p] = v_{kp} \sum_{n=0}^{N-1} x[n] v_{np}, \quad (3)$$

where v_{kp} denotes the p -th component of the k -th CDFT eigenvector. Like the DFT, zero-padding may be used to generate a transform whose size ($P \times P$) is greater than the length of the input vector (N).

3 Chirp Model

The chirp signals considered in this paper are formulated as:

$$s_i[n] = A_i \exp(j(\alpha_i m^2 + \omega_i n + \phi_i)) \quad (4)$$

$$z[n] = \sum_{i=1}^P s_i[n] + w[n] \quad (5)$$

$$m = n - \frac{N-1}{2} \quad 0 \leq n \leq N-1$$

where $w[n]$ is additive Gaussian white noise with a standard deviation of σ . The unknown parameters, $\theta_i =$

$[A_i, \alpha_i, \omega_i, \phi_i]^T$, are the amplitude, chirp rate, center frequency, and phase respectively. This paper focuses on estimating the chirp rate and center frequency. The amplitude and phase are assumed to be unknown for the purpose of Cramér-Rao lower bound derivation, but were fixed at 1 and 0 respectively for all simulations. The algorithms were written assuming they were unknown and did not take advantage of their fixed value.

Many of the applications of chirps measure real-valued signals. The techniques presented here can be applied directly to real valued signals with only minor variation (there will be two 180° -symmetric peaks per chirp instead of one). Alternately, the signal can be converted to complex form using a Hilbert transform. Preliminary investigation suggests that these methods have lower error when operating on complex signals, so that is the focus of this paper.

4 Subspace Decomposition

This paper builds on the work of [3] to use subspace decomposition techniques to improve DFrFT chirp rate estimation in the presence of noise. Subspace decomposition techniques have proven to be very successful at sinusoidal frequency estimation. They are more robust to noise, and more accurate than simple Fourier Transform peak detection. The ideal approach would be to apply a 2D subspace decomposition (such as developed in [9]) however, this is quite computationally expensive. Instead, horizontal and vertical projections of the MA-CDFrFT were calculated us-

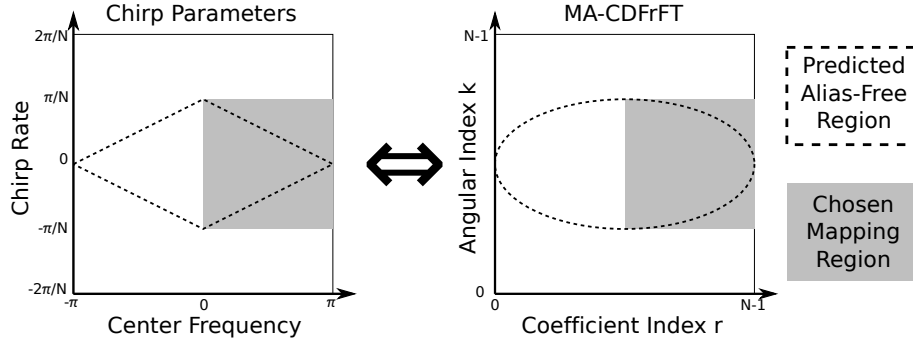


Figure 2: Expected Valid Mapping Region: A mapping must be established between (a subset of) the field of chirp parameters, and the MA-CDFrFT transform grid. Outside of the dashed region, the instantaneous frequency of chirps will be greater than π , and the sampling theorem is not satisfied. We thus only generated a mapping for the region shown in grey.

ing a p-norm:

$$\mathbf{x}_\alpha = \text{FFT}^{-1} \|\mathbf{X}\|_p^{\text{row}} = \text{FFT}^{-1} \left(\sum_{k=K_0}^{K_f} |X_k[\bullet]|^p \right)^{\frac{1}{p}} \quad (6)$$

$$\mathbf{x}_\omega = \text{FFT}^{-1} \|\mathbf{X}\|_p^{\text{col}} = \text{FFT}^{-1} \left(\sum_{r=R_0}^{R_f} |X_\bullet[r]|^p \right)^{\frac{1}{p}} \quad (7)$$

The strong peaks in the projections (as seen in Fig. 1a) result in strong frequency content in \mathbf{x}_α and \mathbf{x}_ω , thus turning the chirp parameter estimation problem into two frequency estimation problems, which were solved using subspace decomposition methods.

We investigated using various p-norms to project the magnitude of the MA-CDFrFT onto the horizontal and vertical axes. The result of our preliminary investigation was that the 1-norm (absolute sum) tended to bury the peak. The 2-norm was not appropriate, as the FrFT respects Parseval's Law at each angle of the transform. The Inf-norm (vector maximum) had a lower noise-floor at high SNR, but was still fairly susceptible to noise at low SNR levels. Compared to the Inf-norm the 3-norm had a higher noise-floor at high SNR, but less at low SNR. It was considered to be a good compromise and was selected for detailed evaluation in this paper.

The subspace decomposition process begins by performing eigenvalue decomposition on estimated covariance matrices \mathbf{R}_α and \mathbf{R}_ω of size $(C + M) \times N$, where C is the number of chirps expected in the signal. The eigenvectors corresponding to the C largest eigenvalues were then selected to form the signal subspace, leaving M eigenvectors to form the noise subspace.

Next the pseudo-spectra was calculated using the MUSIC and Minimum-Norm methods :

$$P_{\text{MUSIC}} = \frac{1}{\sum_{k=1}^M |\text{FFT}(v_k)|^2}, \quad (8)$$

$$P_{\text{MIN-NORM}} = \frac{1}{|\text{FFT}(\mathbf{V}\mathbf{V}^T \mathbf{1})|}, \quad (9)$$

where v_k is the k -th eigenvector of \mathbf{V} , sorted in ascending order, and $\mathbf{V}^T \mathbf{1}$ denotes a column vector containing the first element of each eigenvector. These two algorithms were selected for evaluation in this paper as they showed the most promise of all the subspace decomposition methods examined in [3]. It should be noted that the FFT can be computed for any power-two length R , not just the length of the eigenvectors, which are of length P . Increasing this size scales as $O(R \log R)$, and is much less expensive than increasing the size of the MA-CDFrFT at $O(P^2 \log P)$.

Finally, these pseudo-spectra were searched for the largest C peaks. The indices that were found using this approach differ from the direct 2D peak detection, due primarily to contributions of other terms in the projection. Thus a different peak to parameter mapping function was generated for each method.

5 Generating Peak-to-Parameter Mapping

An empirically determined approximate mapping between the peak location in the MA-CDFrFT and the chirp parameters is given in [2]:

$$\alpha \approx \frac{2}{N} \tan \left(\frac{2\pi}{N} r - \frac{\pi}{2} \right) - \frac{1.41}{N} \left(\frac{2\pi}{N} r - \frac{\pi}{2} \right) \quad (10)$$

$$\omega \approx \left(\frac{2\pi}{N} k - \pi \right) + 0.85 \left(\frac{2\pi}{N} r - \frac{\pi}{2} \right)^3 \quad (11)$$

This approximate mapping has significant error, however, so for this paper we will provide a method to pre-calculate the mapping at each location in the transform to more than adequate accuracy for 2D peak detection.

To calculate the mappings, we generated $N \times N$ sample chirp functions with evenly spaced center frequencies ranging from 0 to π and chirp rates ranging from $-\pi/N$ to π/N . These ranges were chosen because when $|\alpha|(N-1) + |\omega| > \pi$ the instantaneous frequency of the chirp will exceed the range $|\Omega_i[n]| < \pi$, resulting in an

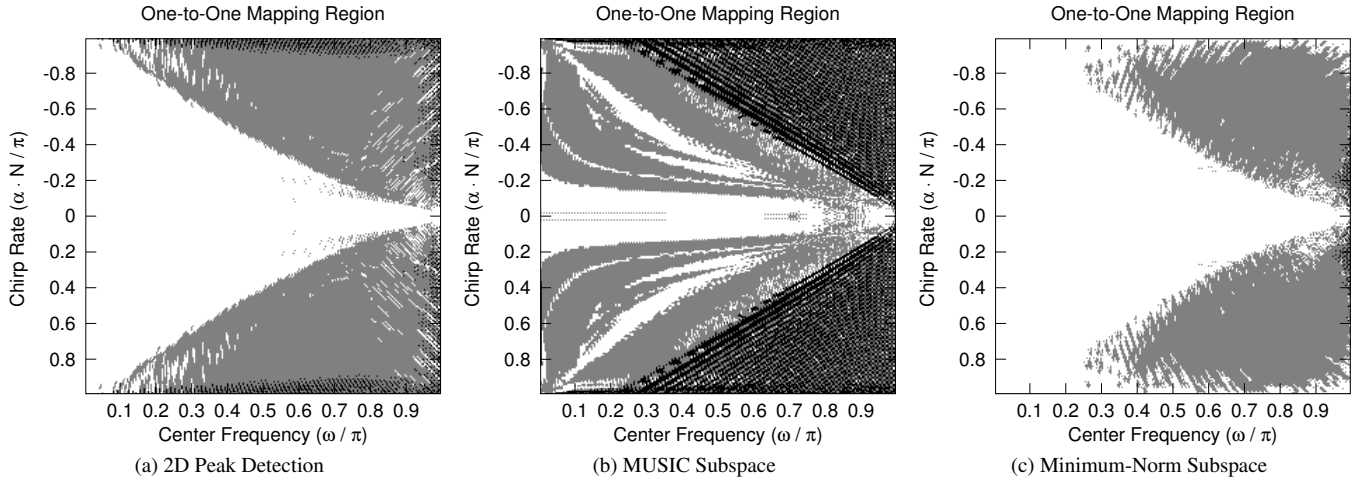


Figure 3: Actual Valid Mapping Regions: The mapping is one-to-one in the white areas, was not one-to-one in the grey regions, and had no solution in the black regions. The area shown in the images correspond to the grey region shown in Fig. 2(a). The mappings were generated for chirps of length $N = 256$, with a transform size of $N \times N$ and have a resolution of $N \times N$. For the subspace decomposition methods, 50 noise eigenvectors were used.

effect similar to aliasing (area shown in Fig. 2). The MA-CDFrFT is thus not well behaved outside of this range and we would not expect the mapping to be one-to-one in this region.

For the direct 2D peak detection methods, we could have restricted our mapping to chirps within this region. However, for the projection-subspace approach, since the horizontal and vertical coordinates of the peak are found independently, it is possible to obtain an estimate for the peak outside of the valid region. To calculate the error in this estimate it is necessary to have a mapping for a full rectangular region.

We then used the above peak detection methods to determine the peak indices for each of the sample chirps. The peaks in the MA-CDFrFT occurred roughly between $k = N/4$ to $3N/4$ and $r = N/2$ to N . Thus the $N \times N$ samples in chirp parameter space mapped to about $N/2 \times N/2$ indices, for an average over-sampling of 4 times. To invert this mapping we started by grouping all the chirp parameters that mapped to a specific index location. This location was mapped to the centroid of the chirp parameters that mapped to it. If these parameters did not form a connected region in the chirp parameter grid, then that location of the mapping was flagged as being not one-to-one. Finally, if no chirps mapped to an index pair, then the chirp parameters for that location were determined by linear interpolation of the surrounding points or 4th-order polynomial extrapolation of the entire mapping depending on whether the location was on the interior or exterior of the mapping.

Figure 3 depicts the regions in the α - ω plane where the mappings were valid. Using the basic 2D peak detection method, the region of the MA-CDFrFT with a valid mapping was a slightly smaller subset of the valid region which the sampling theorem alone would suggest. There were even a few spots on the interior of the trian-

gle that were not one-to-one. Using the Minimum-Norm projection-subspace method actually expanded the valid area, while the MUSIC projection-subspace method made it much worse. Inspecting slices of the MUSIC mapping showed that the non-one-to-one areas were mostly the result of dithering between index steps, and not large discontinuities.

These mappings can be used as a simple lookup table when the resolution of the estimator matches that of the mapping, and this was done with the 2D peak detection. In the case of the subspace decomposition, however, we chose to use an output size R greater than the transform/mapping size P . In this case, linear interpolation of the mappings was used to obtain the chirp parameter estimates. It should also be noted that for the subspace techniques, the projections were only performed over the mapping region $k = N/4$ to $3N/4$ and $r = N/2$ to N , not the entire transform.

6 Extension to Multicomponent Chirps

Modifications are needed to extend these methods to signals containing multiple chirps. Particularly, the projection-subspace approach does not provide information to determine which peaks in the horizontal projection match the peaks in the vertical projection, and thus it must be combined with other techniques to obtain a final result. For this paper, we used 2D peak detection to find the peaks, and then used subspace decomposition on the row and column centered at those peaks to refine the result, in what we are calling the *cross-hairs method*.

In addition, to calculate the estimation error, we must define how the estimated parameters are paired with the actual parameters. We selected the permutation which resulted in the lowest total MSE. The chirps selected to evaluate the multicomponent performance are shown in Fig. 4.

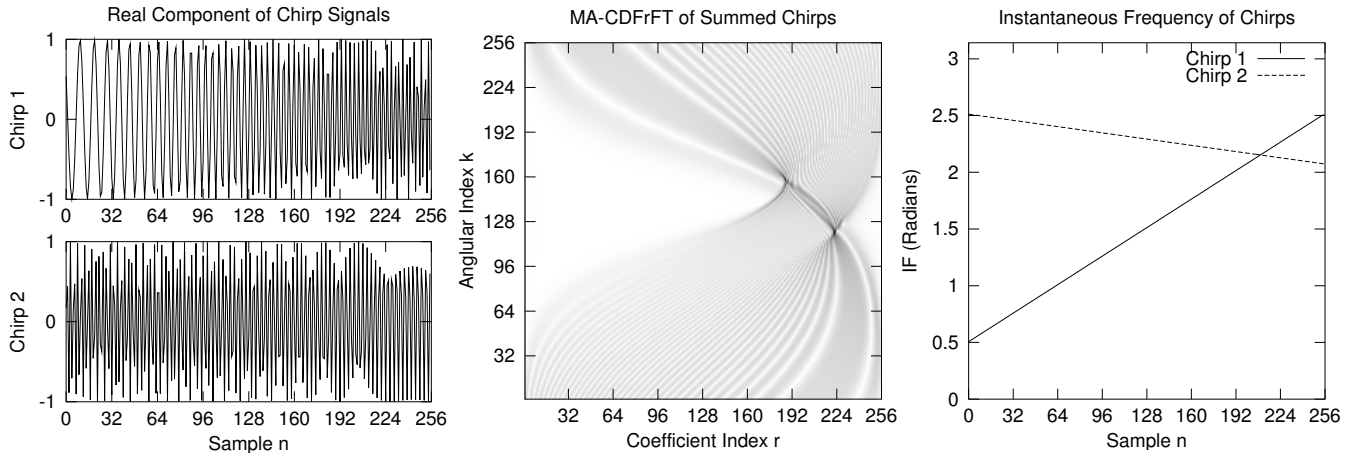


Figure 4: Test Chirp Representations: The chirps used to demonstrate multicomponent estimation performance are shown here in multiple representations to illustrate their characteristics. The first was selected to have relatively high chirp rate of $\alpha_1 = 0.32\pi/N$ and moderate center frequency $\omega_1 = 0.48\pi$. The second is more narrow band with $\alpha_2 = -0.07\pi/N$, and $\omega_2 = 0.73\pi$.

7 Discussion of Results

7.1 Cramér-Rao Lower Bound

To evaluate the performance of these estimators, it is valuable to compare against the theoretical limit provided by Cramér and Rao. This lower bound has been calculated before for chirps [10] [11], but since different papers use slightly different forms of the chirp function, we present our derived results for the specific form used in this paper.

The components of the Fisher information matrix for any signal in complex additive Gaussian white noise is:

$$J_{ij} = \frac{2}{\sigma^2} \sum_{n=0}^{N-1} \left(\frac{\partial \mu_n}{\partial \theta_i} \frac{\partial \mu_n}{\partial \theta_j} + \frac{\partial \nu_n}{\partial \theta_i} \frac{\partial \nu_n}{\partial \theta_j} \right), \quad (12)$$

where $\mu_n = \text{real}(z[n])$ and $\nu_n = \text{imag}(z[n])$, are the expected values of the real and imaginary components of the signal. For the multicomponent case $\Theta = [\theta_1, \theta_2, \dots, \theta_P]^T$ and J will have the form:

$$J = \begin{bmatrix} J_{11} & J_{12} & \cdots & J_{1P} \\ J_{21} & J_{22} & \cdots & J_{2P} \\ \vdots & \vdots & \ddots & \vdots \\ J_{P1} & J_{P2} & \cdots & J_{PP} \end{bmatrix}, \quad (13)$$

composed of the block matrices:

$$J_{ij} = \frac{2}{\sigma^2} \sum_{n=0}^{N-1} \begin{bmatrix} c_{ij}[n] & A_j s_{ij}[n] m^2 \\ -A_i s_{ij}[n] m^2 & A_i A_j c_{ij}[n] m^4 \\ -A_i s_{ij}[n] n & A_i A_j c_{ij}[n] m^2 n \\ -A_i s_{ij}[n] & A_i A_j c_{ij}[n] m^2 \\ A_j s_{ij}[n] n & A_j s_{ij}[n] \\ A_i A_j c_{ij}[n] m^2 n & A_i A_j c_{ij}[n] m^2 \\ A_i A_j c_{ij}[n] n^2 & A_i A_j c_{ij}[n] n \\ A_i A_j c_{ij}[n] n & A_i A_j c_{ij}[n] \end{bmatrix}, \quad (14)$$

where

$$c_{ij}[n] = \cos(\Phi_i[n] - \Phi_j[n]) \quad (15)$$

$$s_{ij}[n] = \sin(\Phi_i[n] - \Phi_j[n]) \quad (16)$$

$$\Phi_i[n] = \alpha_i m^2 + \omega_i n + \phi_i \quad (17)$$

For the case of a single chirp, the inverse of this matrix has a simple closed form, and the Cramér-Rao lower bound is:

$$\text{var}\{\hat{A}\} \geq \frac{\sigma^2}{2N} \quad (18)$$

$$\text{var}\{\hat{\alpha}\} \geq \left(\frac{\sigma}{A}\right)^2 \frac{90}{N(N^2-1)(N^2-2)} \quad (19)$$

$$\text{var}\{\hat{\omega}\} \geq \left(\frac{\sigma}{A}\right)^2 \frac{6}{N(N^2-1)} \quad (20)$$

$$\text{var}\{\hat{\phi}\} \geq \left(\frac{\sigma}{A}\right)^2 \frac{21N^4 - 24N^3 - 66N^2 + 96N - 27}{8N(N^2-1)(N^2-2)} \quad (21)$$

For the multicomponent case, there is no simple closed-form solution, but the matrix inverse can be easily calculated for any specific chirp parameters.

7.2 Single Chirp Performance

With this framework, we proceed to evaluate the performance of the proposed estimation methods using simulated monocomponent chirp signals at various SNR levels. The results are shown in Fig. 5. At low SNR (roughly $< -15\text{dB}$), all the estimation methods return results no better than choosing a point in the mapping range at random. At medium-to-high SNR (roughly $> -5\text{dB}$) the 2D Peaks method is able to determine the center frequency to within the pixel resolution of the transform grid and the average error of the chirp rate is within a few pixels. The subspace methods performed even better in this region because of their higher resolution output.

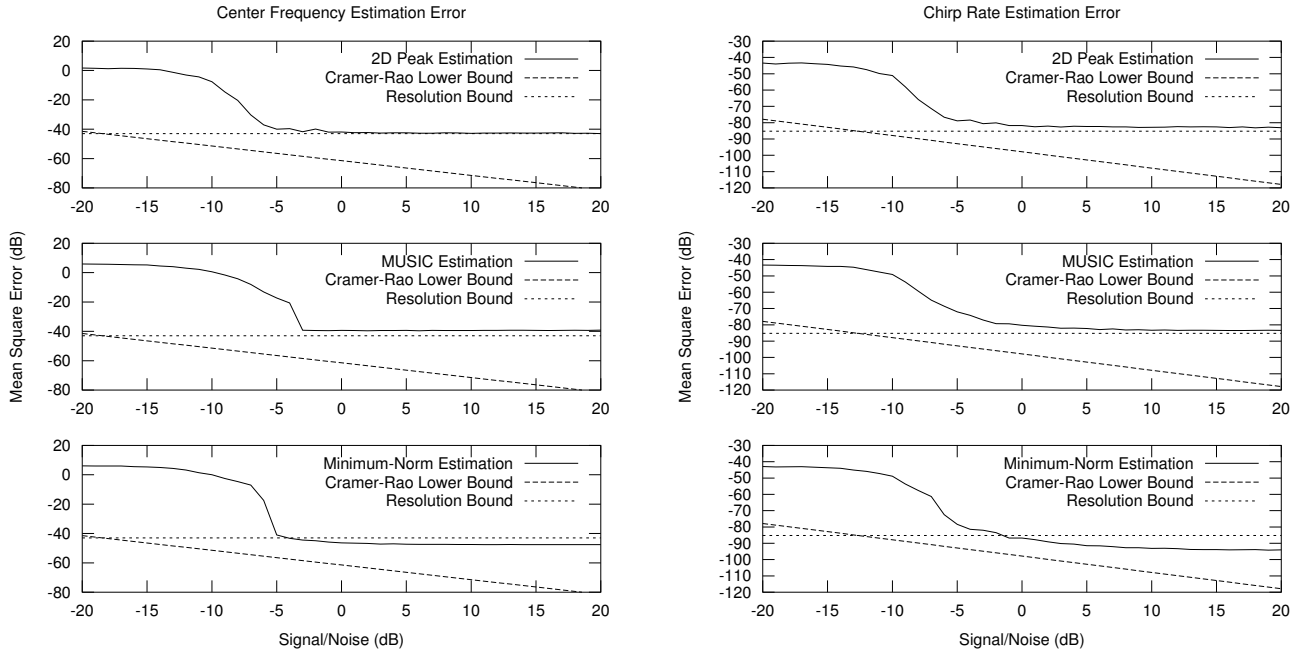


Figure 5: Single Chirp Parameter Estimation Error: The MSE was calculated at each SNR using 1000 chirps of length $N=256$, in the “safe” range of $|\alpha|(N-1)+|\omega| < 0.85\pi$. A transform of size $N \times N$ was used. For the subspace decomposition methods, $M=50$ noise eigenvectors were used, with 3-norm projections as the inputs and an output resolution of $R=2048$.

At medium-low SNR the MSE transitions from these two extremes. Based on preliminary investigation, it appears that this primarily reflects the percentage of time that a noise-generated peak is selected rather than the true peak. In other words, we believe the estimation either provides a relatively accurate answer, or a completely wrong answer in this region.

The subspace decomposition methods actually perform worse in this SNR region. The reason can be seen in Fig. 1. As discussed before, the noise in the signal is distributed throughout the transform, so the peak still stands out above the noise floor in the 2D transform. However, the projections include all the noise from all the rows/columns and the noise floor buries the peak. Furthermore, in the case of multicomponent chirps, the support leading up to each of the peaks sum together to create a high “cross-term floor” even when no noise is present. Using a 3-norm is an improvement over the 1-norm, but still has its limits. Thus the subspace decomposition techniques presented are useful for improving accuracy at low SNR, but not extending the SNR range.

7.3 Multicomponent Performance

The results of 2D-peak detection with two chirps are shown in Fig. 6. These chirps are well separated, and the Cramér-Rao lower bound is only negligibly higher than the bound for a single chirp. Nevertheless, the error begins to increase at higher SNR than for the single-chirp (about 7dB vs -5dB). However, given suitable SNR, they still provide esti-

mations close to the limit of the transform resolution.

The results of the cross-hairs method are shown in Fig. 7. This method performed only slightly better than the 2D peak detection approach. By estimating the chirp rate and center frequency independently, we are treating the estimation as a separable problem, however the approximate mapping given in equation 11 shows that this is not the case. The results are thus consistent with the general principle that joint estimators are more accurate than independent ones.

To better understand what is happening, we generated zoomed-in sections of the mapping near the test chirps, as shown in Fig. 8. As mentioned before, the subspace methods provide super-resolution peak locations, and we must then interpolate the mapping (which is only defined at integer peak locations) to obtain the chirp parameter estimates. For interpolation to be meaningful, adjacent integer peak locations should map to regions of chirp parameters which are also adjacent. Figure 8 shows that this is true in one dimension, but not in both. Therefore, interpolation of the mapping is not valid and instead a separate high-resolution mapping must be generated specifically for the super-resolution subspace method. Generating such a mapping has computational complexity of $O(R^2 P^2 \log P)$, and could not be completed for this paper.

8 Conclusion

In this paper, existing work on multicomponent chirp parameter estimation using the DFrFT and subspace decom-




Figure 6: The multicomponent estimation error using the 2D peak detection method was calculated using the chirps depicted in Fig. 4, averaged over 1000 experiments at each SNR.




Figure 7: The multicomponent estimation error using the minimum-norm cross-hairs method was calculated using the chirps depicted in Fig. 4, averaged over 1000 experiments at each SNR. Subspace decomposition was performed on 1-pixel slices centered on the peaks, with an output resolution of $R=2048$

position was refined to: (a) address peak to parameter mapping issues by explicit evaluation of the relevant peak-to-parameter mapping, and (b) alleviate noise accumulation issues arising from projection operations, using the cross-hair approach. The performance of the refined subspace algorithm was evaluated by comparing it to the CRLB. Simulation results indicate that the proposed approach provides a computationally viable suboptimal alternative to the two-dimensional subspace chirp parameter estimation approach.

References

- [1] X. Jingping, L.-G. Durand and P. Pibarot, Extraction of the aortic and pulmonary components of the second heart sound using a nonlinear transient chirp signal model, *IEEE Trans. on Biomed. Eng.*, 48(3), 2001, 277-283.
- [2] J.G. Vargas-Rubio and B. Santhanam, The centered discrete fractional Fourier transform and linear chirp signals, *Proc. of IEEE DSP/SPE Workshops*, 2004, 163-167.
- [3] B. Santhanam and M. Hayat, On a pseudo-subspace framework for discrete fractional Fourier transform based chirp parameter estimation, *Proc. of IEEE DSP/SPE Workshops*, 2011, 360-363.
- [4] L.B. Almeida, An introduction to the angular Fourier transform, *Proc. of ICASSP*, 1993, 257-260.
- [5] B. Santhanam and J.H. McClellan, The discrete rotational Fourier transform, *IEEE Trans. on Sig. Proc.*, 44(4), 1996, 994-998.
- [6] A. Koc, H.M. Ozaktas, C. Candan, and M. Alper Kutay, Digital Computation of Linear Canonical Transforms, *IEEE Trans. on Sig. Proc.*, 56(6), 2008, 2383-2394.
- [7] S.-C. Pei, M.-H. Yeh and C.-C. Tseng, Discrete fractional Fourier transform based on orthogonal projections, *IEEE Trans. on Sig. Proc.*, 47(5), 1999, 1335-1348.
- [8] J.G. Vargas-Rubio and B. Santhanam, On the multi-angle centered discrete fractional Fourier transform, *IEEE Sig. Proc. Letters*, 12(4), 2005, 273-276.
- [9] J.W. Odendaal, E. Barnard, and C.W.I. Pistorius, Two-dimensional superresolution radar imaging using the MUSIC algorithm, *IEEE Trans. on Ant. and Propag.*, 42(10), 1994, 1386-1391.
- [10] R.M. Liang and K.S. Arun, Parameter estimation for superimposed chirp signals, *Proc. of ICASSP*, 1992, 273-276.

[11] S. Peleg and B. Porat, The Cramer-Rao lower bound for signals with constant amplitude and polynomial phase, *IEEE Trans. on Sig. Proc.*, 39(3), 1991, 749-752.

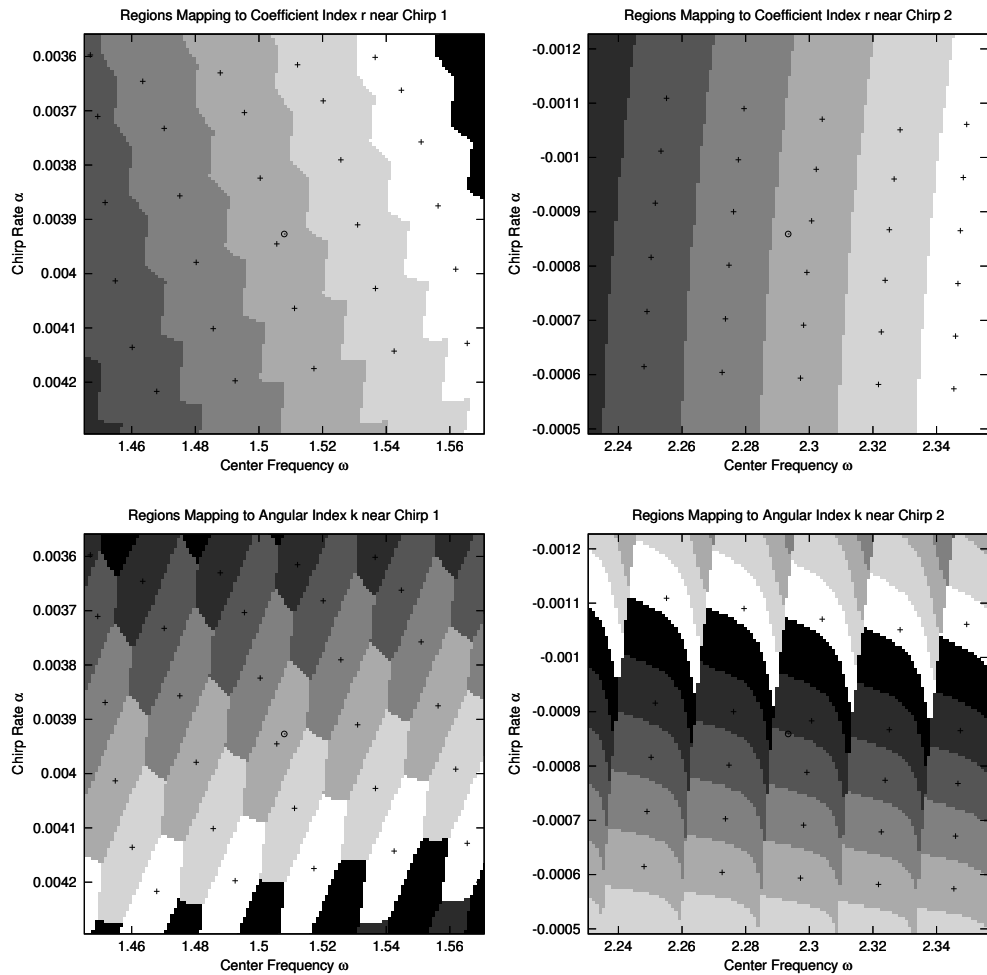


Figure 8: Mapping Bands: These images show segments of the mapping function generated with an oversampling factor of about 400 (20 in each direction), zoomed in near the two test chirps depicted in Fig 4. Each grey level distinguishes the locus of chirp parameters which map to a specific row (or column) in the MA-CDFrFT transform. The cross-hairs indicate the centroid of the chirp parameters which map to a single point in the transform, and the circles indicate the test chirp parameters.